An Operator for Entity Extraction in MapReduce

Ndapandula Nakashole
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA, 15213
ndapa@cs.cmu.edu

arXiv:1512.04973v1 [cs.DB] 15 Dec 2015

ABSTRACT

Dictionary-based entity extraction involves finding mentions of dictionary entities in text. Text mentions are often noisy, containing spurious or missing words. Efficient algorithms for detecting approximate entity mentions follow one of two general techniques. The first approach is to build an *index* on the entities and perform index lookups of document substrings. The second approach recognizes that the number of substrings generated from documents can explode to large numbers, to get around this, they use a *filter* to prune many such substrings which do not match any dictionary entity and then only verify the remaining substrings if they are entity mentions of dictionary entities, by means of a text join. The choice between the *index-based* approach and the *filter & verification-based* approach is a case-to-case decision as the best approach depends on the characteristics of the input entity dictionary, for example frequency of entity mentions. Choosing the right approach for the setting can make a substantial difference in execution time. Making this choice is however non-trivial as there are parameters within each of the approaches that make the space of possible approaches very large. In this paper, we present a cost-based operator for making the choice among execution plans for entity extraction. Since we need to deal with large dictionaries and even larger large datasets, our operator is developed for implementations of MapReduce distributed algorithms.

Categories and Subject Descriptors

H.1.0 [**Information Systems**]: Models and Principles—*General*

Keywords

Approximate Entity Extraction, MapReduce, Text Analytics Optimization

1. INTRODUCTION

Dictionary-based entity mention extraction has wide use in search and semantic web related work. For example, shopping portals annotate text documents with dictionary products to maximize product relevance to user queries. Semantic search detects mentions of entities such as people, organizations and locations in text documents in order to facilitate entity-oriented search.

Determining wether a document mentions a product is a challenging task, particulary becuase entity mentions in Web documents are noisy. Such mentions are rarely exact matches of the dictionary entities, which can be too long for users to write in full everytime they refer to the entity. For example, in product reviews, it is common for reviewers to use short representations of the entities in product catalogs. A mention in a document may miss some of the words of the entity or it may have extra words not found in the entity name in the dictionary, therefore it is crucial for algorithms to detect *approximate mentions* of entities in addition to exact mentions [3, 8, 7]

Finding entity mentions entails finding substrings in a document sequence such that the substring matches an entity in the dictionary. Such *candidate substrings* , are substrings whose words are a full or partial subset of the words of an entity in the dictionary. A naive method to generate *candidate substrings* would be to scan a document, generating all substrings of up to size L where L is the longest dictionary entity and do a dictionary lookup for each of the substrings. Efficient algorithms for this problem follow one of two approaches. The first approach is to build an *index* on the entities and perform index lookups of document substrings [3, 6]. The second approach recognizes that the number of substrings generated from documents can explode to large numbers, to get around this, they use a *filter* to prune many such substrings which do not match any dictionary entity [5] [8] and then only verify the remaining substrings if they are entity mentions of dictionary entities, by means of a text join. If the longest entity in the dictionary has L words. Given a document d, all substrings with length up to L are possible mention candidates. This produces $L \times |d|$ substrings to be looked up. The filter serves to reduce the number of lookups as only substrings that pass the filter are verified if they in fact refer to a dictionary entity or not.

The choice between the *index-based* approach and the *filter & verification-based* approach is a case-to-case decision as the best approach depends on the characteristics of the input entity dictionary, for example frequency of entity mentions. For non-distributed algorithms, the main differentiating factor is that the filter is a more compact structure than the index, in most cases the filter fits in memory whereas the index does not. Performace of the index-based approach is affected by the indexing scheme. Indexing on single words has a different effect to, for example, indexing sets of words that frequently co-occur in dictionary entities. Individual word index posting lists can grow too long which incur long times for merging posting lists. Performance of the filter & verification-based approach is affected by the type of filter used to prune substrings, a basic filter like a prefix-based filter for pruning out prefixes not

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 2011 ACM 978-1-4503-0493-1/11/02 ...$10.00.

likely to match dictionary entities would perform differently from a probabilistic filter such a Latent Signature Hash (LSH) filter.

If we consider the indexing scheme to be a paremeter and the filter to be a paremeter, than the number of available approaches can be large and the choice becomes a challenge. In this paper we propose to add to this space of approximate entity mention extraction methods. Our proposal is based on the fact that while performance of the input entity dictionary at hand, the dictionary can be heterogenous such that it consists of partitions that if each partition is processed seperately by the most suitable method, it can result in cheaper aggregate run times than any of the methods applied to the entire dictionary. This is the hybrid approach. Given the choices of indexing schemes and filters, and the hybrid approach combining index-based and fiilter &verification based methods, the important problem we address is this paper is that of which of the indexing schemes to use, which of the filters to use and wether or not to use a hybrid approach, and if using a hybrid approach where should we partition the entity dictionary. This is an optimization problem over a large space of approaches.

We solve this optimization problem over distributed algorithms, since we need to deal with large dictionaries and even larger large datasets, we develop and optimize over scalable algorithms in the form of MapReduce distributed processing. In a MapReduce setting, network time plays a significant role on job completion time, furthermore, additional MapReduce-specific coordination tasks such as disk-based sorting introduce costs that are not part of single machine algorithms. Additional distributed setting variables make the choice between the *index-based* approach and the *filter & verification-based* approach more challenging. To differentiate between these parallel computation coordination tasks and actual processing time spent on the job, we make a distinction between *work done time* and *job completion time* in the objective functions we define for optimization.

In this paper, we introduce an operator named the *Entity Extraction Join Operator, EE-Join Operator* for optimizing entity extraction in MapReduce. We define the space of approaches for approximate entity mention extraction, and propose a hybrid approach. We develop a cost model for estimating the costs of each of the approaches and efficiently search the space of available options.

In summary we make the following contributions:

- An operator, *EE-Join* for highly scalable and optimized entity extraction in MapReduce, which works with different objective functions, we use two distinct objective functions, the *work done time* and *job completion time*
- A study the space of available approaches for approximate entity mention extraction, and propose an additional approach
- A cost model for estimating execution time for the twor objective functions.
- A means to gather data statistics needed leveraged by the cost model.
- An efficient algorithm for searching the space of available approaches.
- An experimental evaluation of our operator, with entity dictionaries consisting of entities that follow various mention distributions.

The rest of this paper is organized as follows. We next give the semantics of dictionary entity extraction in Section 2. Section 3 describes approximate entity mention algorithms and presents our adaptations of single machine algorithms to their MapReduce counterparts that form the building blocks of the Entity Extraction Join operator. Section 4 describes the costmodel. Section 5 describes the optimization problem over the cost model as the objective function and presents our solution to the optimization problem. We report our experimental results in Section 6. Section 7 is a review of related work on optimization in text analytics and joins in MapReduce. Finally, we conclude in Section 8.

2. SEMANTICS OF DICTIONARY ENTITY EXTRACTION

In this section, we formally define dictionary-based entity extraction, our focus is on approximate mentions where substrings can be partial matches of the dictionary entitites. The similarity function used plays an important role in the semantics of approximate mentions, we therefore define and motivate similarity functions we use.

Approximate dictionary-based entity extraction takes as input a dictionary of entities, ξ and a document collection ζ to output all pairs (e, d), such that $sim(e, d_{s(i)}) >= \gamma$ where $e \in \xi$, $d_{s(i)}$ is a substring in d, $d \in \zeta$ and γ is a given similarity threshold.

To compute pair-wise similarity $sim(e, d_{s(i)})$, different similarity functions can be used depending on the desired semantics. A commonly used similarity measure is *Jaccard similarity* [13, 16], defined as: $JaccSim(e, d_{s(i)}) = \frac{e \cap d_{s(i)}}{e \cup d_{s(i)}}$. Jaccard similarity is a symmetric measure; where asymmetric semantics are desired a different measure is needed. For example, suppose the dictionary contains two entries, *E1:iPhone Charger* and *E2: Apple iPhone 4 Black or White 32G AT&T*. Assuming a document contains a substring *S1: iPhone 4*. Then $JaccSim(E1, S1) = \frac{1}{3}$, but $JaccSim(E2, S1) = \frac{1}{4}$, although *E2* is semantically the better match. Jaccard Similarity gives a lower score to *E2* because a large fraction of its words are missing in *S1*. A measure that reflects the fact that *S1* is fully contained in *E2* would solve this problem. Prior work has defined an asymmetric measure *Jaccard containment* as: $JaccCont(e, d_{s(i)}) = \frac{e \cap d_{s(i)}}{d_{s(i)}}$.

Thus the Jaccard containment of *iPhone 4* in *E1*, $JaccCont(E1, S1) = \frac{1}{3}$, whereas $JaccCont(E2, S1) = 1$.

We use Jaccard containment as the similarity measure. We note that there are two important variations for the *Jaccard containment* measure. The first Jaccard containment variation tolerates missing words in the approximate mention, as in the example above. The second variation tolerates extra tokens in the approximate mention. Clearly, the semantics fo the Jaccard containment variations are different and one may be desirable in settings where the other is not.

DEFINITION 1: Given an entity e and a document substring $d_{s(i)}$, the Jaccard containment of $d_{s(i)}$ in e, allowing for *missing* words in $d_{s(i)}$, is: $JaccCont^{missing}(e, d_{s(i)}) = \frac{e \cap d_{s(i)}}{d_{s(i)}}$. Allowing *extra* words in $d_{s(i)}$, $JaccCont^{extra}(e, d_{s(i)}) = \frac{e \cap d_{s(i)}}{e}$.

Using Jaccard containment and its variations we can leverage an interesting property that enables efficient computation of Jaccard variants. Given a similarity threshold γ, we can compute all substrings $s_{ij} \subseteq s$ such that the total weight of words in s_{ij}, $w(s_{ij})$ is $\geq w \times \gamma$. All these substrings are approximate mentions according to Jaccard containment, we refer to these as the the Jaccard variants of a string. For example, consider the entity, *Apple iPhone 4 32G*. Suppose the token weights are as follows: *Apple:{1}, iPhone:{8}, 4:{2}, 32G:{1}*. For $\gamma = 0.75$, the Jaccard variants is of the entity

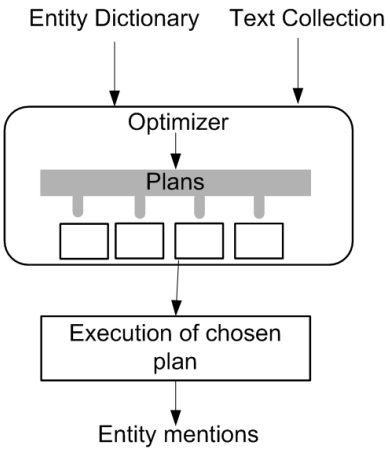


Figure 1: Overview of the EE-Join operator.

are: {Apple iPhone 4}, {iPhone 4}, {iPhone 4 32G}, {Apple iPhone 4 32G}. If we store compute and store all Jaccard variants of dictionary entities, we can perform exact match comparisons between the Jaccard variants of the dictionary entities and the Jaccard variants of potential mentions.

DEFINITION 2: A subsequence $s_{ij} \subseteq s$ is a Jaccard variant of s if $s_{ij} \cap s >= \gamma$. For settings where words are *weighted*, a weighted subsequence $ws_{ij} \subseteq ws$ is a Jaccard variant of ws whose weight is w, if $w(s_{ij} \cap s) >= w \times \gamma$.

Computing the Jaccard variants for the dictionary entities is straightforward, however computing the the Jaccard variants if done naively can explode since every substring in a document is potentially a Jaccard variant of some dictionary entity. All these have to be queried against the variants of the dictionary entities. We avoid generating all possible Jaccard variants as explained later.

3. APPROXIMATE MENTION ALGORITHMS IN MAPREDUCE

Having defined the semantics, we can introduce algorithms for approximate entity mention extraction. For each algorithm we briefly describe the single machine version before explaining how adaption to MapReduce is realized through the use of MapReduce constructs. We then explain the impact of the MapReduce constructs on the performance for each of the algorithms.

3.1 MapReduce SSJoin

Chaudhuri et al. [9] introduced the notion of set similarity join (SSJoin) for identifying similar strings. They observe that typically efficient algorithms for similartity joins use a similarity function chosen to suit the domain and application. The premise of SSJoin is to decouple similarity functions from the implementation of similarity joins, instead they propose an operator as a foundation to implement similarity joins that can adaptively handle a variety of similarity functions. For example, depending on the size of the relations being joined and the availability of indexes, the SSJoin optimizer may choose either index-based plans or merge and hash joins in order to implement the SSJoin operator.

The simplest implementation of the SSJoin on a single machine thus compares every substring s to every entity $e \in \xi$. We adapt the SSJoin algorithm to MapReduce to create a baseline MapReduce algorithm. The mappers generate all substrings s of length l (maximum entity length) for every document $d \in \zeta$. The mappers then generate one or more signatures for each of the substrings. The same mapper functions are applied to the entire dictionary such that for each $e \in \xi$, the mappes generate signatures, applying the same signature generating function. The signature generating function is constructed such that if a substring and a dictionary entity are similar, they will at least have one signature in common. Thus, using signatures as the reduce key, substrings and entities with a signature in common are shuffled to at least one common reducer which computes similarities between substrings. One of the shortcomings of this algoirthm is that it requires a reduce function. This incurs a significant amount of data transfer time required to shuffle the substrings and entities to the reducers. The shuffling cost of this baseline algorithm is exacerbated by the fact that all possible substrings from all documents are generated. The MapReduce algorithm is outlined in Figure 2.

```
function map(item_i, tag)
    if item_i is a document
      List S ← generate all substrings from document (d_i)
    else if item_i is an entity
      List S ← item_i
    for s ∈ S do
      generate signatures from s
      emit signatures

function reduce(sig_i, [item_{1,entity}, item_{12,entity}, item_{31,doc}, ...])
    Hashtable MatchingEntities ← { }
    if item_i is an entity
      add item_{i,entity} to hashtable MatchingEntities
    else
      for item_{i,entity} ∈ MatchingEntities do
        if sim(item_{i,entity}, item_{j,doc}) ≥ γ
          emit(item_{i,entity}, item_{j,doc})
```

Figure 2: MapReduce *SSJoin* baseline

3.2 MapReduce Index on Entities

The index on entities approach creates an index on words of the dictionary entities. It then generates all substrings and queries the index for similar entities. We adapt the index-based approach to entities as follows: First we generate the index on the entities as a seperate MapReduce job. The index is then broadcast to every Mapper node. The dictionary is broadcast to every Mapper node. The type of index used can vary from application domain. Assuming a basic index with inverted lists per word, and then for each query substring q retrieve all lists corresponding to words in q. The union of the lists is the candidates entities that are mentioned by q. Each of the candidates entities are then verified to determine if they are true mentions. The MapReduce algorithm is outlined in Figure 3. Though the index is created in separate MapReduce it is not a significant portion of the execution time as the dictionary is typically much smaller than the document collection. One of the limiting factors of this algorithm is the fact that the index can be large, may not fit in memory. This means the index has to be partitioned into smaller indices which can fit in memory, and the entire corpus has to be processed serveral times, once for every index partition. The MapReduce algorithm is outlined in Figure 3.

The type of index used plays an important role in the performance of the algorithm. We studied three types of indices and their properties.

function map($index_P$, d_i)
 Index $Idx \leftarrow$ load index into memory
 List $S \leftarrow$ generate all substrings from document (d_i)
 for $s \in S$ **do**
 lookup s on the index
 if $sim(e, s) \geq \gamma$
 emit(e, s)

Figure 3: MapReduce *Index-based* lookups

- **Per word index**: This is the basic index, where an inverted list is generated for every word, storing all entities consisting of that word. While single word inverted index can be generated quickly, these lists can grow very large, making the task of list merging expensive.

- **Prefix-index**: A prefix index arranges the words of the entities according to a fixed order, for example, based on decreasing occurrence frequency. During similarity join, we have to generate the prefixes of the substrings and query them against the index to verify that the substring-entity match surpasses a user specified threshold. Like the per word index, the prefix index is quick to generate, its advatange is that it reduces the problem of potentially long inverted lists.

- **Jaccard variant index**: The Jaccard variant index is an index on all the Jaccard variants of all the entities. During similarity join, we have to generate the Jaccard variants of the the substrings. The advantage of the Jaccard Variant index is that it requires no verification, a substring with a Jaccard variant with an inverted list il is an approximate mention of all those entities with the Jaccard variant as a substring. Constructing a Jaccard vairant index is slightly more expensive then other two.

3.3 MapReduce ISHFilter & SSJoin

We have introduced the baseline SSJoin algorithm and the index-based algorithm, both generate all possible substrings and then computing a similarity join between a large set of substrings and the dictionary entities. Perfoming a similarity join between all substrings is a large peformance bottleneck. To overcome this problem we first filter out all substrings that cannot match with any dictionary entity, and only then, perform a set similarity join (SSJoin). We use the The ISHFilter introduced by Chakrabarti et al. [5] to prune a large number of substrings that are obvious non-mentions. The SSJoin algorithm is then applied to remaining substrings with the dictionary entities. The difference between baseline *SSJoin* and *ISHFilter & SSJoin* is that shuffling cost is much lower as number of substrings is substantially reduced by the filter. The MapReduce algorithm is outlined in Figure 4.

So far we have introduced the SSJoin as using a signature scheme to generate sinatures such that substrings and entities with a signature in common are shuffled to atleast one reducer. The signature used upon which the data is shuffled plays a significant role in performance. We studied three signature schemes.

- **Single word signatures**: Using each word as a signature has a lot of skew, becuase some words are very common. This results in high shuffling costs. Since each substring and entity consists of many words, this type of signature results in duplicate work at the reducers.

function map($item_i$, tag)
 if $item_i$ is a document
 List $U_S \leftarrow$ generate all substrings from document (d_i)
 List $S \leftarrow$ apply *ISHFilter* to all substrings U_S
 else if $item_i$ is an entity
 List $S \leftarrow item_i$
 for $s \in S$ **do**
 generate signatures from ss
 emit signatures

function reduce(sig_i, $[item_{1,entity}, item_{12,entity}, item_{31,doc}, ...]$)
 Hashtable $MatchingEntities \leftarrow$ { }
 if $item_i$ is an entity
 add $item_{i,entity}$ to hashtable $MatchingEntities$
 else
 for $item_{i,entity} \in MatchingEntities$ **do**
 if $sim(item_{i,entity}, item_{j,doc}) \geq \gamma$
 emit$(item_{i,entity}, item_{j,doc})$

Figure 4: MapReduce *ISHFilter & SSJoin*

- **Prefix signatures**: The prefix signature uses prefixes as signatures. While quick to generate, the prefix signatures are susceptible to skew, causing shuffling costs to be skewed to a few nodes. The prefix signature requires verification at the reducers.

- **Locality-Sensitive Hashing (LSH) signatures**: is an algorithm for solving approximate or exact similarity. It is probabilistic in the sense that it uses a hash on the the input so that similar items are mapped into the same group with high probability. Like Prefix signatures, LSH signatures require verification.

- **Jaccard variant signatures** : Jaccard variants as signatures has the advantage that it reduces data skew, while also not requiring verification at the reducers.

3.4 MapReduce Index on Documents

The fourth approach for approximate entity mention extraction on MapReduce is to create an index on the entire document collection. The input to the mappers are partitions of the document index, the dictionary of entities is broadcast to every mapper. Each dictionary entity $e \in \xi$ is treated as a query which is posed on the index. Each mapper searches its part of the document index for all entity queries. The complete list of mentions is the union of mentions found by each mapper. When the dictionary of entities does not fit in memory, the algorithm makes multiple passes over the document index. Constructing an index on the entire corpus is an expensive operation, this is unlike the approach that constructs the index on the dictionary of entities becuase the dictionary of entities is usually orders of maginitude smaller than the document collection. Since the index on the document collection is usually not available upfront, and constructing it is expensive, we do not persue this approach further in the rest of the study.

3.5 Operator Algorithms

We have aldready eliminated the *Index on Documents* algorithm from the algorithms considered for the EEJoin operator. We further eliminate the *MapReduce SSJoin* algorithm due to its limitation with generating all substrings. Instead we keep the optimized version of SSJoin, the *MapReduce ISHFilter & SSJoin*. We also

keep the *MapReduce Index on Entities* but instead of generating all substrings from documents, we use the filter to effectively have the *MapReduce ISHFilter & Index on Entities* approach. These two algorithms provide a rich set of options as both the allow tuning of the signatures used by the SSJoin and the type of index used by the entity indexing algorithm.

4. COST MODELS

Having described the algorithms for the EEJoin operator, we now present the cost-model used to estimate performance of the each of the algorithms. Using a cost model, we can automatically determine which of the algorithms has the best performance for a given input dictionary ξ and document collection ζ. We consider two objective functions, one for the total work done and another for the job completion time.

The job completion time of the *Index on Entities* approach is made up of two main componets. The first is the substring lookup time denoted by C_{lookup} in Definition 3. The total lookup time is equally distributed among the mappers due to the MapReduce load balancers for the mapper, thus total lookup time is $\left(\frac{|C|}{|M|} \times C_{lookup}\right)$. The second is the number of iterations made over all the substrings due to th entity dictionary not fitting in memory, denoted by $\frac{|E|}{M_e}$. *explain estimation of |C| and |M| from data statistics.*

DEFINITION 3: The cost of the index approach , for job completion time, is defined as follows: $Cost^{index} = \left(\frac{|C|}{|M|} \times C_{lookup}\right) \times \frac{|E|}{M_e}$ Where $|C|$ is the number of candidate substrings from all documents in the dataset, $|M|$ is the number of mappers, $|E|$ is the dictionary size and $|M_e|$ is the memory budget for the index. Thus $\frac{|E|}{|M_e|}$ is the total number of passes made over the data.

The job completion time of the *ISHFilter & SSJoin* approach consists of three main components. The first is the cost of generating signatures over $\frac{|C|}{|M|} \times C_{sig(c)}$ this cost depends on the type of signature used. The second is the cost of shuffling the signatures over the network, $(C_{suffle(sig)}$. This cost depends on the number of signatures per candidate. The third is the cost of verifying the candidates of a signature, $C_{verifysig}$ as shown in Definition 4.

DEFINITION 4: The cost of the filter & ssjoin approach, for job completion time, is defined as: $Cost^{ishf\&ssj} = \frac{|C|}{|M|} \times C_{sig(c)} + |Sig|\left(C_{suffle(sig)} + C_{verifysig}\right)$. Where $C_{sig(c)}$ is the cost of generating signatures, $C_{suffle(sig)}$ is the cost of shuffling signatures and $C_{verifysig}$ is the cost of verifying candidates of a signature.

5. OPTIMIZATION

We optimize over a large plan space of algorithms. The plan space is made up of two core algorithms. However, each of the algorithms can be instantiated with several signature schemes. Furthermore, the dictionary is partitioned such that a fraction of the entities is processed by one of the core approaches and the rest are processed by another core approach. Furthermore, any combinations of the different signature schemes form a possible hybrid approach, thus creating a large space of possible plans.

5.1 Plan space

The *EE-operator* optimizes entity extraction by partitioning entities into mention frequency categories. The intuition is that each of the algorithms performs better for entities of certain mention frequencies than the other approach. Therefore, for a given input dictionary and text collection, a fraction of the entities is processed by the *index-based* approach, for some signature scheme, and the remaining fraction is processed by the *filter & verification-based* approach, for some signature scheme, not necessarily the same as the one used by the *index-based* approach.

Therefore, a plan for the *EE-operator* is a combination of the *index-based* approach and *filter & verification-based* approach. Thus the cost of a plan, where α proportion of the entities are processed by the *index-based* approach and β proportion are processed by the *filter & verification-based* approach, is:

$$Cost(plan) = Cost(\alpha, sigX)^{index} + Cost(\beta, sigY)^{ishf\&ssj}$$

When only one of the approaches is used, the approach not used contributes a cost of $zero$ to the plan. The signature schemes are denoted by *SigX* and *SigY* for the *index-based* approach and the *filter & verification-based* approach respectively. Thus searching over the search space is an optimization problem to *minimize* the cost of the plan.

5.2 Searching the Plan Space

In this section we describe the algorithm for searching an optimal plan, ρ, in a large space. The plan is a function of the signature scheme and the entity extraction algorithms.s Suppose that we have three signature schemes *Prefix, Jaccard Variants, and LSH*. We pick any pair of approaches, for example, *index-based* approach using prefix signatures and the *filter & verification-based* using Jaccard Variant signatures and we then seek to determine how to partition the entities. If we do a naive enumeration of the costs at every possible partitioning point, for a dictionary size of N, we do N enumerations. The dictionary is typically large, in the order of millions, thus an exhaustive enumeration would not be practical.

We reduce the number of enumerations by using an efficient search algorithm since the entities are already sorted by occurrence frequency.

1. currentCheapestCost = $min(Cost^{index}, Cost^{ishf\&ssj})$
2. searchRange = 0 - N
3. BinarySearch (searchRange, find new cheapest cost < current cheapest)
4. currentCheapestCost = newCheapestCost
5. searchRange = area bounding newCheapestCost
6. Repeat steps 2-5 over an increasingly narrow search range
7. Until new cheapest is == current cheapest OR $searchRange$ is 0.
8. Emit plan(signature, extraction method)

We repeat the procedure for all pairs of approaches. If we have three signature schemes we have maximum of nine pairs, which is a small constant. For each pair we do binary search over an increasingly small range space, c times, which is another small constant. Therefore the complexity of the search algorithm is $\log N$.

We prove that the algorithm works correctly, that is, its output is the partitioning with the cheapest cost within the joint space signature schemes and entity extraction algorithms. In order to prove the correctness of the algorithm, we need to show that both $Cost^{index}$ and $Cost^{ishf\&ssj}$ are monotonically non-decreasing functions over the sorted space of entities, for any given signature

scheme. Since the entities are sorted based on occurrence frequency, based on the definition 3, $Cost^{index}$, the index cost is only affected by the memory size, how many times the entire collection must be searched is based on the memory capacity and the more the number of the entities the more times we have to search over the collection, thus index cost is monotonically non-decreasing. Based on definition 4, the cost of the ssj is based on cost of shuffling, shuffling cost is highest for the most frequently occurring entities which appear at the beginning since the entities are sorted in descending order of occurrence, thus again ssj cost is non-decreasing. We thus have the following lemma.

LEMMA 1: Given an ordered list of entities in decreasing order of Given a collection of entities ordered in descending order of frequency of occurrence. show this with x, y, like math symbols.

The cost of the filter & ssjoin approach is defined as: $Cost^{ishf\&ssj} = \frac{|C|}{|M|} \times C_{sig(c)} + |Sig| \left(C_{suffle(sig)} + C_{verifysig}\right)$. Where $C_{sig(c)}$ is the cost of generating signatures, $C_{suffle(sig)}$ is the cost of shuffling signatures and $C_{verifysig}$ is the cost of verifying candidates of a signature.

6. RELATED WORK

Chaudhuri et al. [9] eveloped an operator-centric approach for for set-similarity joins. The main difference in our approach is the target for MapReduce algorithms. Sarawagi et al. [6] uses an inverted-index approach to compute set overlap string similarities. The main difference with our approach is that we do not fix on the index-approach but instead allow the operator to make cost-based decisions in choosing the best implementation of approcximate entity mentions.

In terms of optimization for text-centric tasks, [14] introduced an optimizer for choosing between query-based and crawl-based method for various text-analytics tasks in a cots-based way. The optimizer adaptively selects the best execution strategy. The EE operator is specifically targeted for MapReduce implementations.

Afrati and Ullman [2] investigated the problem of efficient joins in MapReduce. Vernica et al. [24] developed algorithms for set similarty joins in MapReduce. Yang, et al. [26] extended MapReduce to Map-Reduce-Merge, in order to allow users to express different join types and algorithms. None of these join approaches propose a cost-model approach for finding the best approach for a given setting.

A number of recent work have studied optimization for MapReduce tasks, though none investigate optimization for approximate mention extraction of entities. optimization. The Manimal system [15] analyses MapReduce programs to do general database-style optimizations. Dittrich, et al [11]. proposed the use of indices in order to improve MapReduce performance for certain tasks. HadoopDB [1] combines relational and MapReduce qualities into one system. However, HadoopDB is designed to be a parallel relational database, it does not optimize MapReduce tasks.

7. REFERENCES

[1] A. Abouzeid, K. Bajda-Pawlikowski, D. J. Abadi, A. Rasin, A. Silberschatz. HadoopDB: An architectural hybrid of mapreduce and dbms technologies for analytical workloads. In *PVLDB, 2(1)*, 2009.

[2] Foto N. Afrati, Jeffrey D. Ullman. Optimizing joins in a map-reduce environment. In *Proceedings of EDBT*, 2010.

[3] Parag Agrawal, Arvind Arasu, Raghav Kaushik. On indexing error-tolerant set containment. In *Proceedings of the ACM SIGMOD*, 2010.

[4] A Arasu, V Ganti, R. Kaushik. Efficient exact set similarity joins. In *Proceedings of VLDB*, 2006.

[5] Kaushik Chakrabarti, Surajit Chaudhuri, Venkatesh Ganti, Dong Xin. An efficient filter for approximate membership checking. In *Proceedings of the ACM SIGMOD*, 2008.

[6] Amit Chandel, P. C. Nagesh, Sunita Sarawagi. Efficient Batch Top-k Search for Dictionary-based Entity Recognition In *Proceedings of ICDE*, 2006.

[7] William W. Cohen, Sunita Sarawagi. Exploiting dictionaries in named entity extraction: combining semi-markov extraction processes and data integration methods. In *Proceedings of KDD*, 2004.

[8] Surajit Chaudhuri, Venkatesh Ganti, Dong Xin. Mining Document Collections to Facilitate Accurate Approximate Entity Matching. In *Proceedings of VLDB*, 2009.

[9] Surajit Chaudhuri, Venkatesh Ganti, Raghav Kaushik. A Primitive Operator for Similarity Joins in Data Cleaning. In *Proceedings of ICDE*, 2006.

[10] J. Dean, S. Ghemawat Mapreduce: Simplified data processing on large clusters. In *Proceedings of the OSDI*, 2004.

[11] Jens Dittrich, Jorge-Arnulfo Quiané-Rui, Alekh Jindal, Yagiz Kargin, Vinay Setty, Jörg Schad. Hadoop++: Making a yellow elephant run like a cheetah, without it even noticing. *PVLDB, 3(1)*, 2010.

[12] A. Gionis, P. Indyk, R. Motwani. Similarity search in high dimensions via hashing. In *Proceedings of VLDB*, 1999.

[13] Hadjieleftheriou M., Li, C. Efficient approximate search on string collections. In *Proceedings of PVLDB*, 2009.

[14] Panagiotis G. Ipeirotis, Eugene Agichtein, Pranay Jain, Luis Gravano. To search or to crawl?: towards a query optimizer for text-centric tasks. In *Proceedings of the ACM SIGMOD*, 2006.

[15] Eaman Jahani and Michael J. Cafarella and Christopher Re. Automatic Optimization for MapReduce Programs *PVLDB, 4(6)*, 2011.

[16] Nick Koudas, Sunita Sarawagi, Divesh Srivastava. Record linkage: similarity measures and algorithms. In *Proceedings of the ACM SIGMOD*, 2006.

[17] N. Mamoulis. Efficient processing of joins on set-valued attributes. In *Proceedings of the ACM SIGMOD*, 2003.

[18] S. Melnik, H. Garcia-Molina Adaptive algorithms for set containment joins. In *ACM Trans. Database Syst., 28*, 2003.

[19] Ndapandula Nakashole, Martin Theobald, Gerhard Weikum Scalable knowledge harvesting with high precision and high recall In *Proceedings of the Forth International Conference on Web Search and Web Data Mining, WSDM,* 2011.

[20] N. Nakashole, T. Tylenda, G. Weikum Fine-grained Semantic Typing of Emerging Entities In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics, ACL,* 2013.

[21] N. Nakashole, G. Weikum Real-time population of knowledge bases: opportunities and challenges. In *Proceedings of the Joint Workshop on Automatic Knowledge Base Construction and Web-scale Knowledge Extraction*, AKBC, 2012

[22] N. Nakashole, M. Sozio, F. M. Suchanek, and M. Theobald Query-Time Reasoning in Uncertain RDF Knowledge Bases with Soft and Hard Rules. In *Proceedings of the Second International Workshop on Searching and Integrating New Web Data Sources, VLDS*, 2012

[23] S. Sarawagi, A. Kirpal. Efficient set joins on similarity predicates. In *Proceedings of the ACM SIGMOD*, 2004.

[24] Rares Vernica, Michael J. Carey, Chen Li. Efficient parallel set-similarity joins using MapReduce. In *Proceedings of the ACM SIGMOD*, 2010.

[25] Gerhard Weikum, Johannes Hoffart, Ndapandula Nakashole, Marc Spaniol, Fabian M. Suchanek, Mohamed Amir Yosef Big Data Methods for Computational Linguistics In *IEEE Data Eng. Bull. 35(3)*, 2012

[26] Hung-chih Yang, Ali Dasdan, Ruey-Lung Hsiao, D. Stott Parker. Map-reduce-merge: simplified relational data processing on large clusters. In *Proceedings of the ACM SIGMOD*, 2007.

[27] Matei Zaharia, Andy Konwinski, Anthony D. Joseph, Randy Katz, Ion Stoica. Improving mapreduce performance in heterogeneous environments. In *Proceedings of the OSDI*, 2008.